LAWRENCE LIVERMORE NATIONAL LABORATORY

Science in the National Interest

Bay Area Scientific Computing Day 2006

Technical Talks

David Bindel
University of California Berkeley

"Computer Aided Design of Micro-Electro-Mechanical Systems: Eigenvalues, Energy Losses, and Dick Tracy Watches"

Resonant Micro-Electro-Mechanical Systems (MEMS) can be used as sensors, frequency references, and filters. Surface micro machined resonators currently under development can be integrated with circuitry, so that soon a "cell phone on a chip" may be possible. For these resonators to be useful for such radio applications, engineers need to understand and minimize the amount of damping in the system. In this talk, we discuss the numerical modeling of damping mechanisms for high-frequency resonant MEMS, and describe how the mathematical structure of the numerical models can be used to make fast algorithms for modal analysis and for the construction of reduced-order models. We then describe HiQLab, a simulation package designed to study resonant MEMS. Using HiQLab, we predicted a previously-unknown mode interference phenomenon which causes substantial variation in the quality of certain high-frequency disk resonators. We explain the reason for this phenomenon, and compare our simulated results to measured data.

Jianlin Xia
University of California Berkeley

*Co-authors: Shiy Chandrasekaran – UC Santa Barbara; Ming Gu – UC Berkeley; Mark Hoemmen – UC Berkeley
Xiaoye S. Li – Lawrence Berkeley National Laboratory*

"Superfast Multifrontal Method for Large Structured Linear Systems of Equations"

We develop a super fast direct solver for linear systems arising from certain PDEs such as elliptic equations. During the Gaussian eliminations of the matrices of these linear systems with proper ordering, the fill-ins have this low-rank property: all off-diagonal blocks have small numerical ranks. For problems on regular grids we reveal the above low-rank property by combining the multi-frontal method and the nested dissection ordering of nodes to generate fill-ins. Matrices with this low-rank property can be efficiently approximated with some Hierarchically Semi-Separable (HSS) representations by Chandrasekaran, Gu, etc. These representations are compact when the matrices have small numerical off-diagonal ranks. To efficiently store and handle the HSS matrices some special binary tree structures are used to represent HSS matrices and a set of HSS tree operations are developed. This procedure leads to a super fast multi-frontal method with cost nearly linear to the matrix dimension. Numerical experiments demonstrate the efficiency of the method, especially for large mesh size problems.

Ming Jiang
Lawrence Livermore National Laboratory

"Progressive Techniques for Efficient Processing of Massive Image Models"

Many scientific and government communities are facing a crisis in their inability to process massive volumes of image data. To address this problem, we develop a framework -- ViSUS: Visualization Streams for Ultimate Scalability -- for progressive processing and visualization of large scale distributed data. The breakthrough technology at the heart of ViSUS is a set of algorithms that selectively choose the data that must be uploaded first in order for the images to be immediately useful. Our strategy is to exploit the coupling between time-critical algorithms and progressive multi-resolution data structures to realize an end-to-end optimized flow of data, from the original source, such as remote storage or large scientific simulation, to the rendering hardware. The algorithms that allow ViSUS to run on any computer platform are known as out-of-core, cache-oblivious algorithms. Because the datasets are much larger than the computer main (core) memory, these algorithms optimize the transfer of blocks of data in and out of the computer's memory as they are processed. To demonstrate the efficacy of our framework, we present experimental results from performance analysis using large scale 2D satellite imageries and 3D scientific simulations.

UCRL-ABS-219473-DRAFT

Kurt Stockinger

Lawrence Berkeley National Laboratory

"Highly Efficient Analysis of Large Data Sets with FastBit"

Modern scientific experiments such as Combustion Simulation, Astrophysics or High-Energy Physics often produce large amounts of data in the order of Terabytes or even Petabytes. Each of the data sets usually consists of tens of dimensions where each dimension is a particular attribute such as temperature, velocity etc. Analyzing these data sets efficiently is a challenging task for most existing analysis tools. Even though there has been three decades of intensive research of indexing technology, most algorithms suffer from the "curse of dimensionality" which makes them impractical for high-dimensional search spaces.

We have developed a compressed bitmap indexing technology called FastBit which does not suffer from the curse of dimensionality. Our indexing technology is very efficient for multi-dimensional queries such as "Find all supernova explosions with temperature > T and velocity > V and brightness > B." In this talk we will provide an overview of FastBit and show how we could significantly improve the performance of large-scale scientific data analysis.

Jiawang Nie

University of California Berkeley

Co-authors: James Demmel, Ming Gu – UC Berkeley

"Global Optimization of Rational Functions"

This talk gives recent advances on the global optimization of multivariate rational functions with or without constraints. The Sum Of Squares (SOS) relaxation is proposed, and its convergent properties will be presented. As an example of applications, the nearest GCD problem can be formulated as minimizing a rational function globally.

Josh Griffin

Sandia National Laboratories, Livermore

"A Parallel, Asynchronous Method for Derivative-Free Nonlinear Programming"

The need for derivative-free optimization software commonly arises in real-world engineering problems where function evaluations can be computationally expensive and noisy. Evaluating the objective and constraint function may involve running cumbersome simulation codes with a run time measured in hours. Attempts to compute derivatives can prove futile; analytical derivatives typically are unavailable and noise limits the accuracy of numerical approximations. In this context, pattern search methods are often advantageous; they are well suited for parallel computing and can continue to optimize despite a lack of differentiability in the underlying model. This talk will focus on a specific derivative free algorithm that uses a pattern-search based approach capable of handling both linear and nonlinear constraints. In order to enhance computational efficiency, function evaluations are performed asynchronously in parallel.

Joshua Schrier

Lawrence Berkeley National Laboratory

"The Charge Patching Density Functional Approach for Semiconductor Nanostructures"

First principles density functional calculations typically involve finding a self-consistent solution to the Kohn-Sham equations, scaling with the cube of system size, which quickly become infeasible for studying large systems, such as semiconductor nanocrystals. As an alternative, an approximate ab initio potential may be constructed by patching together local charge motifs determined from self-consistent calculations on small prototype systems, and the eigenvalues determined using the folded spectrum method for only the few band-edge states of interest. In this talk, I will give an overview of this "charge patching" method, as developed by L.-W. Wang and coworkers. As a demonstration, I will then describe calculations of CdS/CdSe/CdS colloidal quantum dot quantum well nanocrystals, a new class of nanostructures in which the quantum confinement effects may be tuned by modification of the core, well, and shell thicknesses. Using the charge patching method, it is feasible to chart the effect of these parameters on the wave function and optical properties, in structures ranging from thousands to ten-thousands of atoms.

Wendy Doyle
Sandia National Laboratories, Livermore

"FCLib: A Library for Building Data Analysis and Data Discovery Tools"

FCLib is a data analysis toolkit constructed to meet the needs of data discovery in large scale spatio-temporal data. The toolkit is a C library of building blocks that can be assembled into data analyses. Our goals were to build a toolkit which is easy to use, is applicable to a wide variety of science domains, supports feature-based analysis, and minimizes low-level processing. This talk will present the design of a data model and interface that best supports these goals, as well as a few of the custom analyses that have been built with the library.

Rich Vuduc
Lawrence Livermore National Laboratory

"The Optimized Sparse Kernel Interface"

The Optimized Sparse Kernel Interface (OSKI) is a freely available collection of low-level primitives that provide automatically tuned computational kernels on sparse matrices for use by solver libraries and applications. OSKI hides the complexity of choosing data structure and code transformations that yield the best performance for a given matrix, machine, and workload of kernel operations. The interface supports legacy applications, exposes the steps and costs of tuning, and allows for user inspection and control of the tuning process. This talk highlights recent work on OSKI, including distributed memory extensions based on PETSc wrappers around OSKI.

OSKI is inspired by a rich Bay Area research tradition in profiling, memory hierarchy optimizations, and automatic tuning, and is part of on-going research by the Berkeley Benchmarking and Optimization (BeBOP) group.

UCRL-PRES-219175
